


CONNEXIONS



The Interoperability Report

November 1992

Volume 6, No. 11

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

The ROAD to a new IP.....	2
Generic Ultimate Protocol....	11
Source Demand Routing Protocol.....	14
Announcements.....	25
Book Review.....	31
Letter to the Editor.....	31

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.
E-mail: connexions@interop.com.

Copyright © 1992 by Interop Company.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the ConneXions logo are registered trademarks of Interop Company.

ISSN 0894-5926

From the Editor

As has been pointed out many times before in this journal: the Internet is growing, and growing fast! The growth is causing depletion of the Internet 32-bit address space, and stressing the fundamental architecture. Two issues, intimately related, and key to the continued success of the Internet are *routing* and *addressing*. In this edition, we take a snapshot of some current efforts that all aim to “save” the global Internet.

Our first article, “The ROAD to a new IP,” explores the nature of the address space limitation and the efforts underway to move from the current Internet Protocol (IP) version 4 to a new IP version 7. In particular, the article considers the challenges of balancing competing concerns, the selection process that is being pursued, and the nature of the routing-, addressing- and header format-related issues. It contains a summary of the major proposals that are before the Internet community. The article is by David Crocker of The Branch Office.

When considering architectural changes and protocol upgrades, it is always tempting to try to solve every problem and limitation, once and for all. This is exactly the approach of the *Generic Ultimate Protocol* (GUP). GUP is a long-term replacement for any protocol used in computer-communications, and is generic enough, so that with respect to the *ISO Reference Model* it may be positioned at *any* layer of the model, or even embrace multiple layers. In the latter case, the only restriction imposed by GUP is that the layers shall be contiguous (an unfortunate consequence of the problems associated with non-contiguous address masks in IP). Research on GUP is being conducted by a concerned member of the Internet community, and we asked the member to give us an overview of this protocol.

Connectivity in the Internet is provided by a network layer protocol (currently mostly IP). The routing of these internet packets over the facilities of a single organization is controlled by an *intra-domain* routing protocol (RIP, OSPF, IGRP, IS-IS, and so on). Our final article addresses *inter-domain* routing which controls the routing between and across facilities belonging to different organizations. The *Source Demand Routing Protocol* (SDRP)—which is a component of the *Unified Approach to Inter-Domain Routing*—is described by Deborah Estrin, USC; Tony Li, Cisco Systems; and Yakov Rekhter, IBM Corporation.

With INTEROP 92 Fall at the end of October, the November issue had to be produced “before the fact,” hence you will have to wait another month or so for a report from the show.

The ROAD to a New IP

by David H. Crocker, The Branch Office

Introduction

The *Transmission Control Protocol* (TCP) and the *Internet Protocol* (IP) have been in full operational use since 1983. There are tens of thousands of IP networks and perhaps millions of users. By any reasonable measure, IP counts as a massive success. Unfortunately, the phenomenal number of users is stressing the design limits of IP: the 32-bit IP host address space simply is too small for long-term growth—although it was more than ample when chosen, 15 years ago. This article explores the nature of the limitation and the efforts to move from the current IP version 4 to a new IP version 7. (Versions numbers 5 and 6 have been used elsewhere.)

Disclaimer #1: This topic has recently engendered significant debating and politicking within the Internet community, including a crisis of confidence that has shaken the core of the *Internet Engineering Task Force* (IETF), this summer. This article attempts to present the technical issues independent of the political and emotional concerns, hoping to aid the reader in considering the nature of the tradeoffs offered by various proposals. It is essential that the choice be made by an informed Internet community, since the choice will not be easy or obvious, but it *will* affect the future of the Internet. General online discussion about this topic is conducted on the mailing list:

`big-internet@munari.oz.au`

but the reader is reminded to subscribe by sending to `big-internet-request`.

Disclaimer #2: The author of this article is an author of one of the proposals. Disclaimer #1 notwithstanding, the author will attempt to avoid the temptation to indulge his certain bias towards his own (excellent) proposal.

Is there really a problem?

An address space of 32 bits allows reference to billions of hosts. (Well, actually host *interfaces*.) Since current consumption of IP addresses is only in the realm of one or a few million, it might appear that there is no pressing concern.

In fact, there are two different problems. One is an immediate crisis and the other is likely to become one within a few years. The first is the size of the information base that must be maintained within IP routers, for making data-forwarding decisions. The latter is simple exhaustion of the IP address space.

Internet growth is approximately 100% every 12 months. While growth in North America has slowed, growth in Europe is explosive and the Pacific Rim is starting a similar curve. All of this is within the usual markets of business and technical users. If new markets open, such as use of IP in consumer products, a significant fraction of the world's population become candidates for IP addresses.

Note that IP, with subnet addressing, divides an address into:

Network : Sub-network : Host

and uses the Class A/B/C mechanism, mixed with address masks, to vary where the boundaries occur between these sub-fields. All of this fits within 32 bits and network numbers are assigned sequentially, with no relationship between any two network numbers, producing a "flat" address space. No matter what games are played within those bits, there is a limit to the maximum number of addresses that can be assigned.

Routing table size

In the Internet backbone, routers are required to know about all of the networks that can be reached. These routers do not have to know about the different sub-networks or hosts, but they do need to know about each and every network in the entire Internet. This is due to the “flat” address space of IP network numbering.

The fact that two networks are topological neighbors, and that packets to the two may travel most of the Internet using the same path, is of no benefit when constructing routing tables. The two network numbers are unrelated, requiring calculation of routes for each of them. The computational cost of calculating these paths separately has become a serious burden. So a way needs to be found to aggregate table entries. This requires a fundamental change in the nature of IP network addresses, from the current, flat style, to one that has more useful structure. The usual assumption is that a hierarchical scheme will be most appropriate.

Address space

While there is no topological information in the encoding of IP network numbers, there is a scheme for distinguishing “large” networks from smaller ones, via the Class A, B, and C mechanism. The most popular addresses are Class B, since they permit reasonably large networks and there is a substantial number (16,383) of such network numbers available. Very few Class A addresses are possible and Class C addresses are useful only for the smallest networks. Hence, the first wall that IP will hit is an exhaustion of Class B addresses [9]. While there may be a way to get around this wall, the second wall, assignment of all possible IP addresses, eventually will be encountered.

There is debate about the imminence of these two walls. However, it appears that modifications to the assignment of Class B addresses, and modifications to backbone router protocols, will allow some deferral of the first concern. As to the latter, enough bits need to be added to allow for reasonable administration and for a scheme which provides some meaningful assistance in relating topological neighbors. For example, a datagram from Japan destined for Stanford University and another from Japan destined for Berkeley ought to be able to travel most of the same path along the Internet, and most of the routers along the way ought to need one—not two—entries in their routing tables. This is only possible if some portion of the two addresses is the same.

**As long as we've got
this thing open...**

IP works well, but improvement always is possible. One line of thought is that this forced change to addressing provides an opportunity to repair or improve other aspects of IP. For example, the space available for IP options is sometimes viewed as too constrained. Similarly, the mechanism for specifying constraints upon handling (*Type of Service*) has never been viewed as adequate and has not seen significant use.

The remainder of this article considers the challenges of balancing competing concerns, the selection process that is being pursued, and the nature of the routing-, addressing- and header format-related issues. It ends with a very brief summary of the major proposals that are before the Internet community. The article is frankly cursory in discussing most of the issues. The current debate ranges wide and deep in considering fundamental aspects of the IP infrastructure. This topic has engendered a great deal of activity and a great many messages, Internet Drafts and RFCs. This article attempts to summarize the issues and proposed solutions, but the reader is encouraged to join the relevant mailing lists and read the relevant documents.

The ROAD to a New IP (*continued*)

Trading immediacy for innovation

Opinions about the timing of total address space exhaustion vary between 2 years and 20 years. However, a plausible, near-term estimate is 5 years. This means that the Internet has until 1997 to develop, test and install a new addressing scheme. Conservative project management thinking would attend to the difficulties of fielding an entirely new technology in such a large community, and would seek to have the new scheme fully deployed in 3–4 years. This leaves very little time to develop and test a solution.

As those who create and distribute products and services well know, the basic development of a capability is often the smallest part of an effort. Testing, manufacturing, training and installation often consume considerable time and resources. By any reasonable measure, it will take 2–3 years to deploy a solution. This leaves us with only 1–2 years to develop, test and stabilize a solution. A choice needs to be made immediately.

This creates a difficult pressure between immediacy and innovation. Some of the alternatives being considered sound quite appealing, but have minimal or fluid specifications, no base of operational experience, and may represent significant differences from current IP experience. The challenge for these alternatives is to demonstrate that the degree of benefit that will accrue from choosing them is sufficiently great to justify the risk of delaying their deployment. Or else, to convince the community that no extra time is required. The challenge for the “simpler” and more conservative alternatives is to convince the community that they offer fundamental safety and cost-effectiveness.

The biggest danger is one of “creeping feature-ism” with more and more requirements being added to the project. It is easy to succumb to this, since we do not get many opportunities to change the infrastructure. However, each additional requirement makes the total task more complicated and risky. At a minimum, it virtually guarantees delays before the total package of changes can be deployed. Technically, the only clear requirement is to enhance IP’s routing and addressing structure. Everything else is beyond the immediate crisis.

Previous efforts

Concern about IP address exhaustion and routing table size explosion has created a sense of crisis within the IETF community. Almost two years ago, a special effort, called the ROAD (*ROUTing and ADDRESSing*) group was formed to consider solutions. It gravitated towards one option, but did not see quick adoption of its recommendation. But time passed and urgency grew. There has been pressure to select a solution immediately, without extensive exploration and development of options. The *Internet Engineering Steering Group* (IESG) divided the concerns into short-term, mid-term and long-term. Class-B exhaustion and routing table size explosion fall into the first category. IP address space exhaustion falls into the mid-term timeframe. The IESG feels that other issues of general enhancement to IP, such as quality of service, security/authentication, mobility, resource allocation, accounting, and high packet rates can be deferred for “long term” consideration.

There is reasonable consensus that the proposal called *Classless Inter-Domain Routing* (CIDR) [2] will be adequate for the near-term, by modifying usage of current IP addresses. However, some conservative members of the community advise against relying entirely upon this one option and suggest that the “mid-term” option be developed with all due speed, in case CIDR proves inadequate. This would suggest that deployment needs to start during 1994!

However, the IESG felt that none of the options being discussed this past spring was sufficiently well specified to allow an adequate analysis of its capabilities. So, the IESG has called for a review at the November, 1992 IETF meeting, with analysis according to a published set of criteria developed by the IESG using community input [3].

The contenders will make presentations at the IETF meeting and the IESG will later issue its recommendation. With luck, a considerable degree of community preference will have developed by that time. Otherwise, the IESG decision may suffer inadequate community support.

Evaluation criteria

The IESG list divides into the following categories:

- *Changes Required:* What is the basic technical work that must be done, to modify IP-related software to support a proposed alternative? This includes direct modification or replacement of the IP module, itself, in hosts and routers, but also includes concern for changes to directory, network management and security services. In general, it is expected that support software will need to be modified, as will various operations and administration procedures. It should be noted that any change, no matter how small, requires that a new address format be supported. This may well be the most significant impact, and it is unavoidable. Note that protocol modules above IP, such as TCP and UDP, need to be able to pass the larger addresses to the IP module, as do user applications. In fact some applications, such as FTP and NFS, currently use IP addresses in their own protocols and will need changing.
- *Implementation Experience:* Simply put, what is the empirical evidence that supports the viability and appropriateness of the alternative?
- *Large Internet Support:* A goal of supporting 10^9 networks and 10^{10} hosts is cited. A proposed alternative needs to describe how its addressing structure will support these numbers and what effect it will have upon routing architectures and tables. An essential concern is the way in which addresses will be administered. If IP addresses are to be sufficient for truly global communications, how will each user obtain an address?
- *Performance Impact:* IP has demonstrated a remarkable robustness for application in increasingly high-speed networks. There is, then, concern about the performance cost of any changes in IP. A proposed alternative will need to explore this issue carefully.
- *Support for Unchanged IP Hosts:* Changes cause incompatibilities. Even if all systems eventually move to the new scheme, there will be a transition period. Related experience suggests that such a period will be extended, possibly on the order of 10 or more years. Hence, there is a concern for the interoperation between systems using the new scheme and systems continuing to use the current IP infrastructure. A proposed alternative needs to discuss its support for "late adopters."
- *Impact on Installed Base of Users:* Amidst the wide-ranging concerns for technical factors, it is easy to miss the likely impact of this change upon the *people* who work with IP technology, ranging from developers and network administrators, to customer support and sales staff. They represent a massive investment in training and expertise. The extent to which an alternative affects that training needs to be discussed.

The ROAD to a New IP (*continued*)

- *Deployment Plan:* Since systems will not convert to a new scheme instantaneously, a proposal needs to detail the methods by which the Internet and its constituents will convert to its use.
- *Future Evolution:* To what extent does an alternative support continued evolution of the IP fabric, into the arena of long-term issues mentioned above?

Design considerations

Routing protocols, and the theories of route calculation, remain a specialized topic with relatively few contributors. There has been significant effort in this area, recently, with the development of OSPF, IS-IS, IDPR, IDRP and BGP. Happily, the current crisis does not seem to require major changes to these new protocols. The key requirement to facilitate large-scale routing is that addresses of Internet neighbors be related in a manner which allows reducing the number of table entries, for distant routers. This is generally agreed to require some sort of addressing hierarchy, such that the hierarchy relates to the “dominant” topological hierarchy. A hierarchy is a tree-structured orientation, yet networks permit “mesh” attachments between any two nodes. Hence, networks usually are not organized as strict hierarchies. However political, economic, and management constraints do tend to cause network interconnections to follow a hierarchy. In the United States, for example, users tend to attach to their organization’s backbone, which in turn tends to attach to inter-organization providers. These may also subdivide into regional and long-haul carriers.

In addition to the routing-related constraint upon design of the new address space, global administration and end-system uniqueness are requirements. The new addresses must be globally unique, as are the current IP addresses. There also is some debate about the distinction between addressing an end-system machine (or process) versus the current style of addressing the network interface of an end-system.

Global administration requires a distributed basis for dividing the space, so that different places can assign unique addresses. Experience with administration of the *Domain Name System* (DNS) suggests that hierarchical delegation of assignment authority works quite well.

Any of the schemes that require more than 32 bits for addressing forces a change to the header format. Proposals range from simply modifying the current IP header to accommodate the additional bits, up to a complete re-working of the header, according to more modern views of efficiency and modularity.

Styles of addressing

A number of approaches to large-scale addressing are being discussed:

- *Association-based:* Curiously, one approach to solving this problem suggests that end systems continue to use 32-bit addresses, but that the network should consider them to refer to “associations” or end-system pairs, like a “connection ID” in a virtual circuit protocol. At any given moment, it is unlikely that there will 4 billion IP-based “discussions” going on, so that the 32-bit address space would be large enough for uniqueness in identifying simultaneous conversations. This approach suggests retaining the current address field, but using it for such association IDs, with routers performing address translation of an ID into a series of routing decisions. This presumes some mechanism for establishing a temporary relationship between an ID and a route.

- *ID-based*: Somewhat similar to Association-based, this uses long-term, globally-unique IDs which specify individual end-systems. IDs for neighboring end-systems may be entirely unrelated. Hence, these IDs are a form of end-system name, rather than address. Addressing information, used to develop a routing sequence, would be derived dynamically. This is felt to be particularly friendly in supporting mobile hosts, but there also is a question about placing a “name” into every IP datagram along with still-necessary addressing information.
- *Provider-based*: In the realm of classic global addressing, provider-based addresses would identify an end-system in terms of its attachment (or, more precisely, in terms of the end-system’s network’s attachment) to a given network service provider. For administrative ease, provider identification probably would be subdivided according to the country in which the provider is present. This leaves open the issue of referencing providers that are multi-national. The major criticism to provider-based addressing is that user networks would be required to change their addresses whenever they changed providers. There is concern that this might reduce competition.
- *Geography-based*: Also offered as a classic approach, geographic addresses are globally unique, but specify the end-system (network) strictly in terms of its geographic location, on the theory that a geographic hierarchy is a reasonable approximation of the global Internet’s major topology. Originally proposed by Steve Deering, of Xerox PARC who called them *city codes*, the major concern about geographic addressing is determination of the final provider to which the target end-system (network) is attached. The current proposal calls for inter-connection facilities, called *Metropolitan Internet eXchanges* (MIX) for the “last hop” hand-off to the target provider.
- *Source-routed*: Proposed by Dave Clark of MIT, a type of addressing which uses *route fragments* would specify a sequence of addresses. The concatenated sequence would be a kind of source route, but with large granularity, rather than specifying each hop along the way. This is spiritually related to the current Loose Source Route IP option.

Proposals

This article has discussed the nature of the impending and current problems, the process that will be used to evaluate proposals, and the types of technical choices that seem to be plausible. The remainder of this article *briefly* summarizes the major proposals that have been offered.

It is remarkable that every single one of these alternatives appears to be entirely reasonable, according to a legitimate set of criteria. They have competent specifications and appear to be technically viable, on their own. For all that, the proposals are quite different, since they attend to different concerns.

Hence, the Internet community is facing an extremely difficult decision. It cannot simply choose based upon the credibility of the people who are proponents or upon “political” concerns. The community needs to determine what factors are most significant to it. Once it does that, the technical choice is likely be straightforward.

The ROAD to a New IP (*continued*)

The proposals are:

EIP *Extended IP* (EIP) was proposed by Zhen Wang, of University College London, and simply creates a new IP option which specifies the additional addressing bits that are needed [7]. There has been some debate about the performance impact of IP options, in routers, and in general, this proposal has not been the subject of sustained consideration within the IETF, though it represents the smallest imaginable change to the current IP format.

IPAE *IP Address Encapsulation* (IPAE) has been proposed by Bob Hinden, of Sun Microsystems, and Dave Crocker, of The Branch Office. It retains the current IP format, but defines a mini-layer above IP and below the transport layer (TCP or UDP) [4].

This mini-layer contains the new, larger global addresses for the source and the destination. Each *commonwealth* has its own 32-bit IP address space, within which the current 32-bit IP address behavior is retained. When the Internet runs out of 32-bit addresses, however, only those systems which have converted to full IPAE format will be able to communicate both within their commonwealth and to other commonwealths. IPAE addresses are likely to be proposed as a 12-octet hybrid of provider-based and geographic-based form with classic IP addresses in the last 4 octets. Discussion by the IPAE working group is on the mailing list:

`ip-encaps@sunroof.eng.sun.com`

SIP *Simple IP* (SIP) is a recent idea from Steve Deering. It retains IP, but makes it simpler. It removes those IP fields that seem to be of little benefit, and retains those that clearly are needed. It proposes an 8-octet extended address, along geography/provider lines. Transition would fall into the category of “dual stack” in that an end-system and intervening routers would have to support old IP and new SIP, in order to talk with all hosts.

CLNP *Simple CLNP* is the gist of the recommendation from the original ROAD group. This proposal would replace IP with ISO’s *Connectionless Network Layer Protocol*, CLNP. It would require use of other lower-layer OSI protocols, but TCP, UDP, and the rest of the upper layers of the Internet protocols would be retained. The presumed strength of this proposal is its use of an international standard and the existence of some installed base. However, there is no operational TCP over CLNP and the conversion effort is not well understood. Classic 20-octet NSAP addresses were proposed.

TUBA *TCP/UDP with Bigger Addresses* (TUBA) is being renamed to *TCP/UDP on CLNP Addressed Networks* (TUCAN) in order to place a reference to CLNP in its name. Written technical work has been provided by Ross Callon, of Digital Equipment Corporation, and it proposes a conversion which is derived from CLNP but with changes appropriate for Internet usage [1]. Further details have been written by Dave Piscitello, of Bellcore [5]. The benefits of this approach are the same as for simple CLNP. One of its departures is the use of ID-based addressing. Transition will require a “dual stack” operation, as discussed above. Discussion is on the mailing list:

`tuba@lanl.gov`

PIP The *P' IP* (PIP) has been developed by Paul Tsuchiya, of Bellcore, and is a completely new internetworking protocol [6]. Header fields are divided into independent sections, such as for routing, versus handling. Transition requires dual stack operation. Very active discussion by the PIP working group is on the mailing list:

`pip@thumper.bellcore.com`

Nimrod *New Routing and Addressing* ("Nimrod") is a routing architecture framework being developed by Noel Chiappa. Nimrod concerns itself less with header formats and more with a routing architecture and a derivative addressing structure. An address would be a hierarchical series of fields, derived up from the "bottom" of the topology. Also supported would be end-point identifiers, along the lines of ID-based addressing.

Summary This last section is decidedly *not* titled "Conclusions" because it is not possible to make any, yet. This section attempts a capsule description of the philosophical choices before the Internet community. It should be noted, however, that the specific addressing architecture proposals, with each header format proposal, could be replaced, so that some permutations and combinations may be worth considering. In any event, an undoubtedly biased summary of the options is:

- PIP is the most radical and offers the greatest opportunity for long-term functional enhancements. By virtue of its great changes and lack of any operational experience, it also offers the greatest risk.
- Nimrod is not tied to a specific format proposal, so that it may be possible to add a Nimrod routing and addressing scheme to PIP, TUBA or IPAE, if it can be developed and tested in time.
- TUBA pursues use of an existing protocol which is functionally similar to IP but already contains larger addresses. It is, however, also pursuing changes to those addresses. The primary strength of TUBA is its relationship to a well-documented international standard; its greatest weakness is its differences in detail from IP, beyond the necessary addressing differences.
- IPAE preserves current IP formats and software, imposing incremental changes only on those systems desiring full Internet connectivity. Its strength is the amount of software and training that will be retained. Its weakness is its apparent change to the Internet addressing model, by imposing a routing barrier between commonwealths.
- SIP preserves the gist of the current IP formats which are known to be needed and otherwise only changes the size of addresses. The strength of SIP is its apparent simplicity and familiarity; its weakness is its newness and lack of operational experience.

References IETF rules prohibit the citation of documents contained in the volatile Internet Drafts directory of the Internet Repository, which is replicated in various places around the Internet. Nonetheless, it may be the only location for some of the documents cited here (with inadequate information). On the other hand, the reader may find that RFCs have since been issued for such documents.

The ROAD to a New IP (*continued*)

- [0] Postel, J., "Internet Protocol," RFC 791, September 1981.
- [1] Callon, R., "TCP/UDP over Bigger Addresses (TUBA)," RFC 1347, May 1992.
- [2] Fuller, V., Li, T., Yu, J. & Varadhan, K., "Supernetting: an Address Assignment and Aggregation Strategy," March 1992.
- [3] Gross, P. & Almquist, P., "IESG Deliberations on Routing and Addressing."
- [4] Hinden, R. & Crocker, D., "A Proposal for IP Address Encapsulation (IPAE): A Compatible Version of IP with Large Addresses." To be re-issued as "IP Address Encapsulation (IPAE): A Compatible Version of IP with Large Addresses" and "IPAE Implementation & Transition."
- [5] Piscitello, D., "Use of ISO CLNP in TUBA Environments."
- [6] Tsuchiya, P., "The 'P' Internet Protocol."
- [7] Wang, Z., "EIP: The Extended Internet Protocol a long-term solution to Internet address exhaustion."
- [8] Lottor, M., "Internet Growth (1981–1991)," RFC 1296, January, 1992.
- [9] Solensky, F., "The Growing Internet," *ConneXions*, Volume 6, No. 5, May 1992, pp 46–48.
- [10] Marine, A., "How Did We Get 727,000 hosts?" *ConneXions*, Volume 6, No. 5, May 1992, pp 49–51.
- [11] "Information processing systems—Telecommunications and Information Exchange between systems—Protocol for Providing the Connectionless-mode Network Service," ISO 8473, March 1987.
- [12] Hagens, R., "Components of OSI: CLNP or A Day in the life of Ivan CLNPacket," *ConneXions*, Volume 3, No. 10, October 1989.
- [13] *ConneXions*, Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."
- [14] *ConneXions*, Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."
- [15] *Interconnections: Bridges and Routers*, by Radia Perlman, Addison-Wesley, ISBN 0-201-56332-0, 1992.

DAVID H. CROCKER is a principal at The Branch Office, a consultancy specializing in strategic market planning and technical system architectures for communications products and services. Mr. Crocker has participated in the development of internetworking capabilities since 1972, first as part of the ARPANET research community and more recently in the commercial sector. He wrote the current Internet standard for electronic mail header formats (RFC 822), and was a director and principal architect for MCI Mail. Mr. Crocker currently serves as the IETF Area Director for Standards Management, attempting to facilitate the community's development of technical specifications. His recent technical efforts include a proposal for accommodating larger IP addresses and for enhancing the Internet's electronic mail transport (SMTP). E-mail: dcrocker@Mordor.Stanford.EDU.

A Generic Ultimate Protocol (GUP)

by A Concerned Member of the Internet Community

Overview

A *Generic Ultimate Protocol*, GUP, is a long term replacement for any protocol(s) used in computer communications. GUP is generic enough, so that with respect to the *ISO Reference Model* it may be positioned at any layer of the model, or it may even embrace multiple layers. In the latter case the only restriction imposed by GUP is that the layers shall be contiguous (this is just an unfortunate consequence of the problems associated with non-contiguous address masks in IP).

GUP packet format

A *GUP Protocol Data Unit* (GUPPDU) consists of an unlimited sequence of triplets of the form <Type, Length, Value>.

Due to its extreme generality, it is not expected that there will ever be a need in any foreseeable (or even unforeseeable) future to have a new version of GUP. Thus the version field used by some of the current protocols is viewed as completely unjustifiable for GUP (it would do nothing, but introduce an extra byte or two of the overhead).

As a side comment, it is actually expected that GUP would make all other protocols obsolete. Thus future generations would be able to find the concept of version number only in some history books. Similar rationale applies to the reason why the length field is absent.

Using GUP for computer- communications

Computers exchange information by packaging it into GUPPDUs that are transmitted as a sequence of electrons (or photons in case of fiber optics). The protocol is so generic, that it would be able to adapt with no changes to any conceivable and imaginable transmission technology. For example, advances in the elementary particle physics may result in the ability to use particles other than electrons or photons (e.g. quarks) to carry GUPPDUs. No changes to the protocol would be needed.

To achieve maximum flexibility, the protocol places no semantics on the individual triplets. Interpretation of these triplets, if needed, is based on the "consenting adults" principle. That, of course, implies a certain level of maturity of the computers using GUP. Procedures for certifying the necessary minimum maturity level is completely outside the scope of this document.

One of the problems we observed with the current protocols is that in certain cases the headers of these protocols (headers are the places that carry control information) are incapable of carrying all the information the user may ever desire. Typical examples are, DoD IP (RFC 791) which is incapable of carrying DARPA contract numbers in its header, and CLNP (ISO 8473) which is incapable of carrying RSA encrypted credit card numbers in its header.

By placing no upper bound on the number of triplets composing a single GUPPDU, and by eliminating any requirements to define semantics associated with individual triplets, GUP demonstrates an unthinkable advance in this area. Moreover, to achieve unprecedented flexibility and extendibility, the protocol places no assumptions on whether there will even be a need to interpret individual triplets. In fact, one possible use of GUP is to carry unlimited number of placeholders for any imaginable (or even unimaginable) future use by any possible (or impossible) computer or any other entity (see below on using GUP in conjunction with GUTs).

Generic Ultimate Protocol (*continued*)

Applicability statement

From the previous description it should be obvious that GUP is capable of solving literally any problem that may be conceived in the mind of any biological entity that is under the impression that it is in its mind. Advances in the area of Artificial Intelligence may expand the scope of the GUP applicability to non-biological entities as well. Of course, it may be possible that GUP is the only thing needed by AI (reader may note potential recursion; procedure for terminating this recursion is completely outside the scope of this document).

Using GUP to solve the Internet scaling problem

Two of the critical problems the Internet is presently trying to deal with are the scaling problem and the IP address exhaustion problem. Either of these problems, by itself, is likely to severely cripple the Internet.

GUP with its unlimited unprecedented and unthinkable flexibility is guaranteed to offer an extremely simple solution to both of these problems, thus killing two birds with a single shot (killing more than two birds with GUP is expected to be an NP-complete problem). Therefore, we suggest that GUP should be immediately adopted simultaneously as a short term, medium term and the long term solution to the above problems.

Truck #4

Work within the ROAD (*Routing and Addressing*) group resulted in the model consisting of three trucks, namely forwarding table explosion, IP address space exhaustion, and emergence of public data carries. However, practical experience clearly demonstrated that the model exhibits a certain degree of shortsightedness. For example, the model utterly failed to predict an unprecedented explosion in the mental energy of the members of the Internet that resulted in an unprecedented explosion of various ideas expressed in the form of SMTP (*Simple Mail Transfer Protocol*) exchanges. This phenomena, which we called "Truck #4," should not be taken lightly.

A superficial analysis may suggest that the direct consequence of the Truck #4 is the injection of an enormous amount of entropy in the system. According to the classical laws of thermodynamics, once an entropy is injected in the system it can not be removed. Thus one may think that the Internet is at the brink of a total disaster caused by Truck #4.

Fortunately, a more careful and scientific analysis shows that this problem may be solved by GUP, augmented by GUT (*Generic Ultimate Translation*). A GUT is a generalization of the NAT (*Network Address Translation*) concept developed within the ROAD group. All the ideas that are presently expressed in the form of SMTP would be carried as <type, length, value> triplets within either single or multiple GUPPDUs. Reduction of entropy is achieved by passing GUPPDUs through strategically placed GUTs. (One possible place for GUTs are NAPs.)

Of course, it is not clear whether it will actually be necessary to employ GUTs (since the protocol is not required to interpret individual triplets). Thus the GUTs will be deployed on demand, as needed. It is expected beyond the reasonable doubts that caching would solve all the GUTs "hot-spot" problems, if any. Further research in this area is unlikely to be needed, but if needed it is outside the scope of this document (but inside the scope of one or more of GUPPDUs).

GUP, combined with GUT, allows us to avoid long and counter-productive discussions on the subject of whose solution is better. GUP+GUT would allow each individual to immediately deploy his/her favorite solution in any place in the Internet. GUP+GUT would provide seamless integration of provably incomparable solutions, thus opening new frontier in the area of formal logic in general, and computer communication in particular.

Future work

Since by definition GUP is the ultimate protocol, no further work on solving any problems will be needed (at least in the area of computer networking). That creates a shift in the Internet paradigm, where instead of working on a solution to a problem people would work on a problem for a solution. The impact of this shift in the Internet paradigm on the volunteer participants is one possible subject for future work.

Further work is needed to study applicability of GUP's entropy reducing capabilities to other areas affected by the increase in the entropy. Of course, GUP should be immediately introduced in all courses on thermodynamics and the fundamental laws of thermodynamics need to be revisited to reflect GUP's phenomena.

Introduction of GUP in the course material may be simplified by using the Internet capabilities. It is expected that such usage would not violate any AUPs.

Security considerations

If GUP is used to embrace multiple layers of the OSI reference model, the embrace may be deadly. On the other hand, it was observed that there is a clear distinction between the reference model and the referenced reality. By following the principle that "Guns don't kill, people do" we may conclude that a deadly embrace in the model will have no impact on reality.

References

- [1] "Information processing systems—Telecommunications and Information Exchange between systems—Protocol for Providing the Connectionless-mode Network Service," ISO 8473, March 1987.
- [2] Postel, J., "Internet Protocol," RFC 791, September 1981.
- [3] Hagens, R., "Components of OSI: CLNP or A Day in the life of Ivan CLNPacket," *ConneXions*, Volume 3, No. 10, October 1989.
- [4] Malamud, C., "The Ultimate File System," *ConneXions*, Volume 5, No. 4, April 1991.
- [5] *ConneXions*, Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."
- [6] *ConneXions*, Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."
- [7] Lottor, M., "Internet Growth (1981–1991)," RFC 1296, January, 1992.
- [8] Solensky, F., "The Growing Internet," *ConneXions*, Volume 6, No. 5, May, 1992.
- [9] Marine, A., "How Did We Get 727,000 hosts?" *ConneXions*, Volume 6, No. 5, May 1992.

A CONCERNED MEMBER OF THE INTERNET COMMUNITY has been involved for many years in the development of computer communication protocols for internetworks. The member may occasionally be seen at IETF meetings or be reached via e-mail as: `ACMotIC@The.Internet.Org`.

Source Demand Routing Protocol:

A Component of The Unified Approach to Inter-Domain Routing

by

Deborah Estrin, USC

Tony Li, Cisco Systems

Yakov Rekhter, IBM Corporation

Introduction

Over the past decade the Internet evolved from a single network connecting a relatively small number of computers to a global infrastructure. The infrastructure is composed of numerous organizations that provide networking resources (e.g. routers and links), as well as an even faster growing population of end-systems. The growth rate we are experiencing today is likely to result in 10^9 to 10^{12} computers interconnected via the Internet before the end of this century. Hand in hand with the growth in the number of computers, the Internet has witnessed growth in the number of organizations whose collaborative efforts provide ubiquitous connectivity between all the computers. Many commercial service providers, such as AlterNet, ANS, EBONE, PSI, SprintLink, AlterNet, CERFnet, etc... (the list is far from being exhaustive and their number is growing rapidly) have now joined government-sponsored facilities (NSFNET, NSI, ESnet) in the provision of the Internet services.

Connectivity in the Internet is provided by a network layer protocol (currently either IP or CLNP). The routing of these internet packets over the facilities of a single organization is controlled by the *intra-domain* routing protocol (e.g., RIP, OSPF, IGRP, IS-IS,...). This article addresses *inter-domain* routing which controls the routing between and across facilities belonging to different organizations.

Requirements

Design objectives for inter-domain routing are driven by two fundamental, sometimes conflicting, requirements:

- The need to provide ubiquitous connectivity for a very large number of computers (10^9 to 10^{12}), and
- The need to support some form of control over resource usage.

Providing ubiquitous connectivity within the Internet, implies that some organizations must share their networking resources with other organizations, and rely on other organizations to do the same. However, because of the costs associated with deploying and managing network resources, and the internal reliance on network availability, organizations must retain control over network usage.

Models for inter-domain routing

Administrations that allow their networking resources to be used by other administrations are called "connectivity providers." Administrations that use the networking resources of other administrations are called "connectivity subscribers." Some administrations are both subscribers and providers. A connectivity provider should be able to select or constrain its connectivity subscribers. Likewise, a connectivity subscriber should be able to select its connectivity providers. The combination of these two control mechanisms is essential to support a global internet composed of multiple organizations.

Growth in the number of connectivity providers comprising the Internet and rapid advances in transmission and switching technologies are likely to yield growth in the range of services offered by providers. While some of the services may be ubiquitous, others will be not be available uniformly across the Internet.

In addition, some services will be offered on a selective basis due to policy concerns; for example, a particular provider may allow only select subscribers to use its networking resources. We use the term *custom routes* for those that support services that are not widely accessible, available, or utilized. In contrast, *generic routes* are ones that support services that are accessible, available, and utilized by a large number of subscribers.

Using the concept of connectivity providers and connectivity subscribers we define two models of interaction between providers and subscribers that must be supported by inter domain routing. We will use the terms providers and subscribers for brevity.

Generic model

In the first model the service offered by a provider is represented by the set of destinations that can be reached via the provider's communication facilities. In other words, once a subscriber establishes direct connectivity to a provider, the subscriber can reach all the destinations offered by the provider. Since the value of any communication service depends upon the scope of reachable destinations, interconnection among providers is encouraged. A typical example of this is the creation of the *Commercial Internet eXchange* (CIX) which interconnects multiple connectivity providers like AlterNet, CERFnet, and PSI. The CIX enables an AlterNet subscriber to have connectivity not only to other subscribers directly connected to AlterNet, but to subscribers connected to all the CIX members, e.g., CERFnet and PSI. In practice the interconnection of various providers creates a mesh of providers and subscribers. We expect that in most cases a provider will calculate and advertise a route to each reachable destination and thereby support a large portion of internet traffic. Since the resulting routes are distributed to all subscribers, we call this the *generic routing model*.

Custom model

A different model may be realized if the service offered by a provider consists solely of information about who may use the provider's resources to access other directly connected providers, and does not involve any information (either direct or indirect) about resources outside of the connectivity provider's control). In this model a subscriber selects all the connectivity providers that may be required to meet the connectivity requirements and computes its own route(s). We call this the *custom routing model*.

Tradeoffs

The two models, generic and custom, present some basic tradeoffs. On the positive side the first model has very nice scaling properties, since a given subscriber has to deal only with the limited number of providers (the number is constrained by the number of the providers to which the subscriber attaches directly.) On the negative side, the choices available to a connectivity subscriber are limited by the choices made by the service providers, just as a consumer's choices are limited by the design choices made by product manufacturers and the buyers in the retail stores. On the other hand, the providers do not make their choices in a vacuum, but based on the needs of their subscribers. Thus, in practical terms, choices offered by a provider reflect the general requirements of its subscribers.

Conceptually, the second model is more flexible with respect to the choices available to the subscriber. This is because the subscriber may have information about more than just directly attached providers and thus may have more choices with respect to how it selects the providers that are necessary to accommodate its desired connectivity. With this model, a subscriber always has substantially more information than what would be available with the first model.

continued on next page

Source Demand Routing Protocol (continued)

To extend the metaphor, in this model the consumer is able to select from all available components and then construct a custom designed product. This model places more responsibilities and demands on the subscriber. As a consumer, the subscriber must be more aware of the alternatives that are available and must carry more information and make its own decisions. The benefit is that the subscriber now experiences fewer constraints, and is better able to obtain the connectivity that it may need.

In practical terms the flexibility promised by the second model is predicated on the ability to utilize the information in a meaningful fashion. The growth of the Internet causes growth in the volume of this information. This makes the promised flexibility more and more problematic, unless there is some means of limiting the amount of information needed by the subscriber to obtain the desired services. This reflects rather poor scaling properties of the second model if it is used as the sole inter-domain routing mechanism.

Unified Architecture

The analysis of the two different models suggests that the best way to accommodate routing in large heterogeneous internets with various types of services and policies is a hybrid model that utilizes the strengths of each model. This combination is the essence of the Unified Architecture. By combining generic and custom routing, the Unified Architecture provides support for inter-domain routing in internets of practically unlimited size, while at the same time providing efficient support for custom routing requirements.

As an example of this unified approach consider the inter-domain topology shown in Figure 1. Most of the domains attached to the connectivity provider W subscribe to the "standard" service provided by W. W, in turn, provides this standard service by subscribing to the service provided by X. However, another domain Y which is also connected to W provides a "special" service (e.g., a new resource-reservation mechanism for real-time applications). Domains B and E attached to W wish to use both the "standard" service provided by W as well as the "special" service that they can obtain by going through Y. This functionality can be accomplished efficiently by using generic routes to support the "standard" service, while B and E use custom routes to access the "special" service.




Figure 1: Illustration of Unified Architecture

Components

In the Unified Architecture, support for two different types of routing is accomplished via two distinct components. Analyses presented in [1,3] concluded that the most suitable technique for supporting generic routing is "path-vector" routing algorithms in combination with clustering techniques, while the most suitable technique for supporting custom routing is source routing. We refer to the first as *Node Routing* (NR), and the second one as *Source Demand Routing* (SDR).

In the case of NR, the existing *Inter-Domain Routing Protocol* (IDRP) and *Border Gateway Protocol* (BGP), provide a suitable architecture [6, 7]. The addition of SDR is an enhancement to the protocol in which custom routing requirements are supported by the SDR component instead of by configuring BGP or IDRP databases with the additional custom routes.

The choice for a particular protocol for the SDR component was more difficult; in part because it represents a greater departure from conventional routing techniques. Over the last few years the ability to support SDR functionality has been explored by the *Inter-Domain Policy Routing (IDPR) Working Group* of the IETF and culminated in an architecture, protocols, and implementation. (Note that IDRP and IDPR are two distinct routing architectures.)

IDRP versus SDRP

Our motivation for designing an alternative to the existing IDPR protocols is to build on the lessons learned in designing and implementing IDPR (with the IDPR working group) and to produce a simpler and more efficient protocol that still provides the source-demand routing capability first addressed by IDPR. The following list outlines the important differences between the *Source Demand Routing Protocol* (SDRP) and IDPR:

- Global distribution (flooding) of dynamic routing information is replaced by on-demand, limited distribution mechanisms.
- Route setup and associated state installation and management is optional.
- Lighter-weight setup mechanism that relies on end-to-end reliability instead of hop-by-hop.
- Employs existing inter-domain routing mechanisms to track adjacent-domain reachability, instead of building an additional protocol (i.e., IDPR's virtual gateway protocol).
- Processing and route computation mechanisms designed to complement a node-routing protocol such as BGP/IDRP.

BGP and IDRP were described in recent *ConneXions* articles [5, 8]. Therefore, the rest of this article focuses on the new SDRP and its interaction with BGP/IDRP.

SDRP overview

The purpose of SDRP (pronounced “*esdrip*,” to go with “*eyedrip*,” which is how IDRP is pronounced) is to support source-initiated selection of inter-domain routes, to complement the intermediate-node selection provided by BGP/IDRP.

The protocol makes minimal assumptions about the distribution and acquisition of routing information needed to compute the SDRP routes. In the next section we describe one possible technique for dealing with partial routing information—the technique of maps.

Routing with maps

The Internet can be modeled as a collection of domains (aka “autonomous systems”) interconnected via inter-domain links. A domain may impose transit policies which specify the restrictions attached to the use of the domain's internal resources. A map is a description of connectivity between domains, possibly augmented with transit policies of the domains on the map. Source routes can be constructed using such maps. While at a sufficiently high level of abstraction, routing with maps may be viewed as the same as routing based on link state databases, a closer examination reveals noticeable differences.

Source Demand Routing Protocol (*continued*)

These differences are:

- A map is not expected to reflect the operational status of links and nodes that provide the inter-domain connectivity (e.g., whether a particular link is up or down at a particular moment in time),
- Maps are not expected to be up to date, and
- Maps are not required to contain complete global information about inter-domain connectivity (e.g., you can have a map of only one particular country, or you can have a map of only research networks, etc.).

(For further discussion of the use of incomplete maps see [2, 3, 4]).

Each domain may acquire its map in a number of ways. The map may be obtained from an out-of-band database, via a different protocol, or by some hybrid combination of these techniques. One example of such a mechanism would be to extract topology information already carried by BGP/IDRP, and also have BGP/IDRP provide information pointing to topology and policy servers that can be queried in the remote domain. Different domains may have different techniques for acquiring maps.

Once the domain has a map, it then constructs one or more source routes based on available maps. Source routes can be constructed well in advance of actual use. To provide redundancy a given domain may construct multiple routes to a given destination. Each domain may have its own algorithm for constructing source routes. Once a route is constructed, forwarding along the route is accomplished using the SDRP mechanisms described in the following sections.

Border Router

A *Border Router* (BR) is a router that forwards packets between distinct domains. Typically, a BR in one domain will connect directly to a BR in the second domain and the pair together effect the inter-domain connection.

Route Servers

We assume that each participating domain has one or more *Route Servers* (RSs) that are capable of constructing domain-level source routes (i.e., a sequence of domain identifiers indicating the sequence of domains through which packets should travel from source to destination). Domain-level source routes do not specify individual routers and links.

In general, RSs may be distributed or centralized. They may be co-resident with hosts or with BRs.

Explicit source routes

The basic mechanism for SDRP packet forwarding is to provide a domain-level source route for each packet.

A source route carried by a packet is expressed in terms of a sequence of domain identifiers. A source route may be either strict or loose. In the former case a packet with the source route carries the exact sequence of domains the packet has to traverse. In the latter case the sequence describes not all, but some of the domains along the path. The choice between strict and loose source route is left to the entity that constructs the route.

Since current versions of IP and CLNP do not provide sufficient functionality to support domain-level source routes, this first realization of SDRP must rely on *encapsulation*.

Encapsulation

Encapsulated packets are carried within the Network Layer protocol that is native to the interior of whatever domain is currently traversed. Encapsulated packets carry as data the packets generated by the end-systems (hosts) within the source domain, as well as the domain-level source route. Therefore, inter-domain packets have the following structure if viewed in transit:

- Header of network currently being traversed (often IP) with the destination field set to the exit BR that is along the path to the next hop in the domain level source route.
- Encapsulating header including domain level source route.
- Original header of network protocol used by the end-system that generated the packet (often IP) with destination field set to the address of the destination end-system(s) (e.g., IP address or NSAP).
- Data carried inside the end-system's packet.

Since the first version of SDRP will be implemented for the Internet, we will assume for the remainder of this article that the outer header formats is IP and we have only to define the encapsulating header, its processing, and how the values are set in the outermost IP packet header (for CLNP networks simply substitute CLNP for IP).

The encapsulating header must contain the following information:

- *Encapsulation type*: to accommodate different types of encapsulation in the future, e.g., when route setup/installation is used instead of explicit source routes in every packet.
- *Probe bit*: if set, the decapsulating router sends a probe reply message to the encapsulating router.
- *Loose/strict domain-level source route flag*.
- *Control packet flag*: indicating if the packet contains control (e.g., error) information or end-system data.
- *Next hop pointer*: to indicate the current position along the source route, i.e., which domain identifier to treat as the next hop.
- *Source route length*: length of the source route.
- *Domain level source route*: sequence of domain identifiers.
- *The address of the BR that originally performed the encapsulation*; this insures that control messages are returned to the BR that performed the encapsulation.
- *Protocol family of innermost packet* (host packet): e.g., IP or CLNP.

Table entries

The originator's BR acts on behalf of end-systems within its domain (in the future some hosts may be equipped to generate the SDRP packets themselves but we do not assume this functionality currently). Initially BRs will be configured with the list of local hosts and destinations (and optionally other information as well) that require source routes. For example, in the Table 1 the first entry indicates that the host with IP address 192.9.200.2 can exchange packets with any host on Network "35" via one of the two routes (each route is expressed as a sequence of autonomous system numbers), listed in priority order.

Source Demand Routing Protocol (continued)

Similarly, the second entry indicates that any internal host on Network “129.34” can exchange packets with any host on Network “129.140” via one of the two routes listed. (The table can include more than just host address information as the index, for example type of service bits that may be encoded in the packet header.)

Host	Destination	Source Routes (precomputed and prioritized)
192.9.200.2	35.*	< 750, 145, 233 >, < 750, 184, 293 >
129.34	129.140.*	< 266, 184, 149 >, < 149 >

Table 1: Example of BR’s configuration

Individual domains determine what bindings, i.e., table entries, are desired/permitted and how they are established. Eventually some domains may have a protocol to establish these bindings dynamically; others will use configuration mechanisms only. For now it is adequate to rely on existing router configuration mechanisms.

The BR creates a header with the indicated source route information and encapsulates the end-system’s original packet. The encapsulated packet is then transmitted within a network packet (in our initial case, IP) with the destination field set to the address of a BR along the path to the next hop domain.

When a border router receives an encapsulated packet, it uses the pointer as an offset into the source route. The destination field in the outermost packet header is updated to reflect the border router that is along the path to the *next* domain in the source route. If the source route is strict, the indicated domain must be adjacent to the local domain, or the packet is dropped and an error message returned. If the source route is loose then the indicated domain may not be adjacent to the local domain and the packet may be routed through one or more intermediate domains before having the next-domain-pointer advanced further. In either case, at the last border router on the route, the encapsulating header is stripped off and the internal packet is sent to the destination end-system.

The only thing that a router needs in order to forward an SDRP packet is a route to the next-hop domain (indicated in the outermost IP header) and routes to its adjacent domains. If the source route is a correct strict route (at the domain level), then this information must be available via BGP/IDRP. If the source route is a loose route, then this information may not be available via BGP/IDRP. In such a case, the router considers it a type of failure and sends a control message back to the source indicating that the required route can not be completed due to the lack of information (see “Control Messages” below).

Routing to domain identifiers

To forward packets along SDRP routes a domain must carry information about other routes to other domains. At a minimum, it is expected that a domain has routes to all of the adjacent domains. Such routes are injected into the domain by its border routers. For example, if a border router A within domain X is connected to a border router B in domain Y, then it is expected that A will advertise a route to Y to all border routers within X, either via BGP or IDRP.

Once a border router detects that its peer in an adjacent domain is unreachable (this detection is outside the scope of SDRP), it shall advertise the route to that domain as unreachable via its node routing protocol.

Control messages

A border router will not forward an SDRP packet if the indicated next hop domain is unreachable due to link or node failure, changes in configured inter-domain topology, or changes in policy.

If such an error is encountered en route, the detecting router responds by generating an error message back to the originating border router. The payload of the error message includes an appropriate error code and possibly portions of the SDRP packet. A border router that has to send back a notification should try to send this notification via a path that doesn't violate transit policies. For example, the border router may use a path available via BGP/IDRP or via SDRP itself. In the latter case we assume that control messages are not subject to transit policies, but that, in general, transit policies should be respected when possible.

Exempting the control packets from transit policies allows the source domain to receive the negative feedback. Any negative feedback is viewed as just an optimization that compliments the positive feedback (via the probe mechanism).

When a failure message is received by the originating border router, the route used by the original offending packet is extracted from the error-message payload. Depending upon the error code, routes are marked or replaced. Details as to the best way to use this dynamic information in the handling of cached routes and cached status information is a subject of ongoing research and initially we will simply invalidate the route and use an alternative.

Probing source routes

When a source route is constructed by the RS it is based on old information, and it is possible that the destination is not actually reachable via this route. Further, it is possible that either an SDRP error message cannot return to the source, or that the source routed traffic is falling into a black hole at some point along the path. To prevent long lived black holes, the encapsulating source router may set the probe bit in the encapsulating header. When the last decapsulating BR sees that this bit is set, it is obligated to return a control message back to the encapsulating BR that originated the source route. This control message contains the original source route and an indication of success. The encapsulating border router may use this positive feedback as an indication that the source route is currently operational and that further traffic may be sent using this route, at which time the bit is turned off. The encapsulating border router may also choose to set the probe bit at intervals to help detect if a route is still viable.

Enforcing transit policies

With SDRP a domain may enforce its transit policies by applying filters to transit traffic. The transit domain may cache both positive and negative information so that it can filter efficiently.

Computing the domain level source routes

A Route Server maintains a map and can receive updates to this information via configuration information, BGP/IDRP, error messages relayed from a border router, or from some future mechanism for limited distribution of dynamic updates. Conceptually, this information may be stored in terms of *route fragments*. A fragment may be as long as the whole path, or as short as a single inter-domain hop.

A fragment has transit policies attached to it. A fragment need not contain transit policies for all source-destination domain pairs. Transit policies associated with a given fragment are the policies applied to the transitive closure of pertinent transit policies of all domains along that fragment.

Source Demand Routing Protocol (*continued*)

If a fragment contains only a single domain, then it contains only the pertinent policies of that domain. For a given route fragment, an RS need not keep transit policies that are irrelevant to the domain the RS belongs to. For example, transit policies for sources outside the domain need not be retained.

A fragment may have an operational status and timestamp attached to it. That status tells whether at the indicated time in the past the fragment was operational or not.

A route computed by BGP/IDRP, or a subset of such a route, may be used as the basis of a fragment, but it will not necessarily include complete policy information. Future versions of SDRP may include additional mechanisms for querying or otherwise obtaining information from the policy servers along the route, as well as for caching more up to date information about recently used SDRP routes or fragments; however these are subjects of ongoing research and are not essential for the initial version of SDRP [2].

In summary, the information present in a route server will not always reflect either current transit policies or current operational topology. However, the scheme provides notification when an attempt is made to either violate transit policies or forward over non-operational segments of the topology.

Initial deployments

In practice, a domain may know a priori what route it wants to use as a primary, secondary, tertiary, etc... Thus, for the initial deployment, it is assumed that each border router within a domain will be configured with source routes, but no dynamically-acquired or computed fragments. These source routes will be applicable on a per destination basis, but apply only to the hosts within the source domain. This requires a border router to have information about addresses internal to the domain. For the initial deployment there is no requirement to dynamically gather information about topology and/or policy, other than by explicit SDRP error messages received when trying to use a source route and via the Probing mechanism.

Status

The SDRP protocol specification and usage documents are being prepared for submission as internet drafts. We hope to have a first prototype implementation by December of 1992.

Future developments

Future developments of the SDRP protocol will include the following:

- More sophisticated and efficient algorithms for computing SDRP routes: One possible algorithm is to extract an initial route from BGP/IDRP. This route will contain the path of domains traversed to the destination, and may be compared against the fragments in the map. If this route does not violate any policies, then this route may be used immediately. In this case, the situation can be optimized by not encapsulating the packets, since they will normally follow the BGP/IDRP path. If the BGP/IDRP path does contain a policy violation, the algorithm might choose to query the topology and policy server immediately prior to the domain which cannot be traversed. From this point, the algorithm might perform a directed depth first search, adding the acquired information to its map. The algorithm may also use other fragments already in its map. Note that the algorithm may include other heuristic information, source specific search preferences, and other out of band information. Once an acceptable route is found, this may be further inserted into the map as its own fragment.

- Route setup/installation option: In some situations the added bytes per packet of an explicit source route may be deemed inefficient. On a more functional level, if the communication involves sending identical packets to multiple destinations and routing requires SDRP and not traditional hop by hop routing (e.g., due to TOS requirements), the source may wish to install a corresponding multicast forwarding tree. It is not practical to include and process this tree on a per packet basis. In this case, installation of the forwarding information may be warranted.
- Limited distribution techniques for dynamic topology information: A source would benefit from more up to date information about the status of inter-domain links that it uses often. Therefore we will distribute dynamic status changes to sources that are currently passing through the affected domain. These sources can be identified by referring to cached forwarding filters or to route or reservation establishment information.
- SDRP relies on a route computation procedure (i.e., an algorithm) to take route fragments, destinations, and type of service parameters as input, and construct complete paths as its output. For the first phase of SDRP it will be adequate to use a simple search algorithm, or even to manually construct the source routes. However, procedures for efficient route computation in this context are a critical area of ongoing research.

Conclusion

By combining NR and SDR routing we propose to support inter-domain routing in internets of practically unlimited size, while at the same time providing efficient support for customized routing requirements.

The development of this architecture and protocols does assume that routing requirements will be diverse and that custom routes will be needed. On the other hand, the architecture and protocols do not depend on assumptions about the particular types of routes demanded or on the distribution of that demand. Routing will adapt naturally over time to changing traffic patterns and new services by shifting computation and installation of particular types of routes between the two components of the architecture.

The proposed architecture combines hop-by-hop path-vector and source-routed protocols, and uses each for that which is best suited: Node Routing uses path vector and multiple, flexible levels of confederations to support efficient routing over commonly used routes; Source Demand Routing uses source routing to route over special routes. In the past, the Internet community has viewed these two as mutually exclusive. The Unified Architecture shows, that to the contrary, they are quite complimentary and it is fortunate that we, as a community, have pursued both approaches in parallel. Together these two approaches combined in the Unified Architecture will flexibly and efficiently support both generic and customized routing in very large and heterogeneous global internets.

Acknowledgments

Hans-Werner Braun (San Diego Supercomputer Center), Scott Brim (Cornell University) and Steve Hotz (University of Southern California) provided valuable feedback on a previous draft.

Source Demand Routing Protocol (*continued*)

References

- [1] Lee Breslau and Deborah Estrin, "Design and Evaluation Of Inter-Domain Policy Routing Protocols," *Journal of Inter-networking*, Volume 2, No. 3, 1991.
- [2] Lee Breslau, Deborah Estrin, Daniel Zappala, and Lixia Zhang, "Exploiting Locality to Provide Adaptive Routing in Real-Time Flows in Global Internets: Abstract," IEEE Workshop on Multimedia Communications, April 1992.
- [3] Deborah Estrin, Yakov Rekhter, and Steve Hotz, "Scalable Inter-Domain Routing Architecture," In Proc. ACM SIGCOMM, 1992.
- [4] Deborah Estrin, Yakov Rekhter, and Steve Hotz, "A Unified Approach to Inter-Domain Routing," RFC 1322, 1992.
- [5] Susan Hares, "Components of OSI: Inter-Domain Routing Protocol (IDRP)," *ConneXions*, Volume 6, No. 5, May 1992
- [6] Kirk Lougheed and Yakov Rekhter, "A Border Gateway Protocol 3 (BGP-3)," RFC 1267, 1991.
- [7] "Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs," ISO/IEC/ JTC1/SC6 DIS10747.
- [8] Yakov Rekhter and Dave Katz, "The Border Gateway Protocol," *ConneXions*, Volume 5, No. 1, January 1991

[Deborah Estrin is reachable c/o Computer Science Department, University of Southern California, Los Angeles, California 90089-0782, e-mail: estrin@usc.edu. Tony Li is reachable c/o Cisco Systems Inc., 1525 O'Brien Drive, Menlo Park CA, 94025, e-mail: tli@cisco.com. Yakov Rekhter is reachable c/o IBM T.J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY, 10598, e-mail: yakov@watson.ibm.com. The work of D. Estrin was supported by the National Science Foundation under contract number NCR-9011279. The work of Y. Rekhter was supported by the Defense Advanced Research Projects Agency, under contract DABT63-91-C-0019. Views and conclusions expressed in this article are not necessarily those of the Defense Advanced Research Projects Agency and National Science Foundation.]

DEBORAH ESTRIN is currently an Associate Professor of Computer Science at the University of Southern California in Los Angeles. She received her Ph.D. (1985) in Computer Science and her M.S. (1982) in Technology Policy, both from MIT. She received her B.S. (1980) in Electrical Engineering Computer Science from U.C. Berkeley. In 1987 Estrin received the National Science Foundation, Presidential Young Investigator Award for her research in network interconnection and security. Her current research interests include: inter-domain routing for global internets; reservation protocols and adaptive routing to support new high-speed, real-time services; and simulation and emulation of large scale networks. She chairs the IAB's Autonomous Networks Research Group and is a participant in the IDPR Working Group of the IETF. Dr. Estrin is a co-Editor of *Internetworking: Research and Experience*, a refereed research-journal published by Wiley.

TONY LI received his B.S. from Harvey Mudd College and Ph.D. from the University of Southern California. He currently slings bits for Cisco Systems, Inc., specializing in IP exterior routing protocols and fixing nasty bugs.

YAKOV REKHTER received his M.S. in Physics from St. Petersburg (former Leningrad) University, Russia (former USSR), his M.S. in Computer Science from New York University, and his Ph.D. in Computer Science from Polytechnic University. Yakov is a manager of High Performance Networking group at the T.J. Watson Research Center, IBM Corporation. He is a chairman of the BGP Working Group of the IETF and an active participant in the ANSI X3S3.3 committee. He is one of the principal designers of the NSFNET Backbone routing architecture and protocols.

Call for Participation

The *USENIX Symposium: UNIX Applications Development* will be held in Toronto, Ontario; Canada, March 29–April 1, 1993. The event is co-sponsored by the USENIX Association and UniForum Canada.

- Purpose** One of the major uses of UNIX today is the support, development, and execution of applications ultimately used in achieving end users' business goals. The current trends in large end-user organizations of downsizing major applications from older mainframes to less expensive, more powerful, and simpler, modern, networked, machines lend UNIX a serious position in the commercial marketplace. Consequently, more and more computing and information systems professionals are encountering UNIX when developing and maintaining applications. The purpose of this symposium is to expose the challenges of building and maintaining applications on UNIX platforms, to discuss solutions and experiences, and to explore existing practice and techniques.
- Format** This symposium will feature papers, invited talks, panel discussions, and tutorials on aspects of designing, building, testing, debugging, and maintaining applications within and for the UNIX environment. There will also be ample opportunity at this symposium to meet your peers and make contact with others with similar interests. This symposium will provide valuable information to designers, programmers, and managers who are planning to port existing applications into the UNIX environment or move development and maintenance teams from proprietary environments to UNIX.
- Suggested topics**
- Graphical User Interfaces:*
 - The X Window System.
 - User Interface Design & Standards.
 - Open Look, Motif, NeWS, and so on.
 - What is a style guide?
 - Importance of consistency and ease of use.
 - Porting Issues:*
 - Issues surrounding the tasks of porting an existing application to UNIX, as well as issues of making UNIX applications portable to other architectures and other platforms.
 - Networking:*
 - Client–Server design issues, etc.
 - Project Management:*
 - Using UNIX tools to support project management.
 - CASE—What, When, Why, Who, How.
 - O/S Issues:*
 - Overcoming limitations set by hardware and operating systems.
 - Security:*
 - The impact of security features.
 - Schemes for maintaining security within an application.
 - Transaction Processing:*
 - Implementing distributed transaction processing for UNIX applications.
 - Fourth Generation Languages:*
 - What advantages and disadvantages do 4GLs have in a UNIX environment?

Call for Participation (*continued*)

Distributed Applications:

- How do you make the best use of existing UNIX functionality (such as e-mail) to build UNIX applications? What are the issues of building and/or using distributed databases?

Object Oriented Programming:

- Productivity, languages, techniques, case studies, etc.

Object Oriented Databases:

- Advantages, etc.

The Corporate Internet:

- High Speed for the Elite, or Connectivity for the Masses?
- ISDN, TCP/IP, OSI, UUCP.
- Governments, privateers, service providers, co-operatives, telcos.
- Philosophy: open road, tollbooths, freeloaders or lifeblood.

Delivering/Installing Applications:

- What's the best way?
- How to prevent piracy, worms, viruses, etc.
- How to do updates effectively and securely.

Testing & Certifying Binary Applications:

- Who does this?
- What does this achieve?
- How long does it take?
- Applications and POSIX.1 Conformance Testing.

Standards:

- ABI/API/ANDF—How, What, Where, When, Why? What are they? How are these standards used? How do they affect applications? What features does each have? What benefits are derived from using each? Where should they be used/followed? When will they be real? How do you keep up with new standards? Why are they necessary?

Submission details

Papers may feature real-life experiences, as well as research topics. Both case-study and technical papers will be accepted. Case studies should describe existing systems and include implementation details and may also include performance data where practical. Submissions must be in the form of extended abstracts (1500–2500 words; 3–5 pages in length). Shorter abstracts might not give the programme committee enough information to judge your work fairly and, in most cases, your submission will be rejected. Longer abstracts and full papers simply cannot be read by the committee in the time available. Feel free to append a full paper to an extended abstract; this is sometimes useful during evaluation. The extended abstract should represent your paper in *short form*. The committee wants to see that you have a real project, that you are familiar with the work in your area, and that you can clearly explain yourself. Please note that presentations are usually scheduled to last 25 minutes. Your presentation should provide an overview of your paper and entice your audience to read it in the proceedings and hopefully follow up on your solution, or take your advice into consideration.

Papers will be judged on technical merit, relevance to the theme, and suitability for presentation. Papers are welcome from software (and hardware) vendors who wish to share their innovative solutions and techniques, but be fore-warned that product marketing will not be tolerated.

Persons interested in participating in panel discussions should contact Greg A. Woods, woods@usenix.org.

Tutorials

Explore topics essential to successful use and development of UNIX and UNIX-like operating systems, The X Window System, networking and interoperability, advanced programming languages, and related areas of interest. The USENIX Association's well-respected tutorial programme offers you introductory and advanced, intensive yet practical tutorials. Courses are presented by skilled teachers who are hands-on experts in their topic areas. In an effort to continue to provide the best possible tutorial slate, USENIX is soliciting proposals for new tutorials. If you are interested in presenting a tutorial, contact the Tutorial Coordinator, Dan Klein, dvk@usenix.org, Tel: +1 412-421-2332.

Invited talks

As part of the technical sessions, a series of invited talks provides introductory and advanced information about a variety of interesting topics, such as using standard UNIX tools and employing specialized applications. We welcome suggestions for topics as well as request proposals for particular Talks. In your proposal, state the main focus, include a brief outline, and be sure to emphasize why your topic is of general interest to our community. The Interim Invited Talks and Panel Coordinator is Greg A. Woods, woods@usenix.org.

BOFs

Birds-of-a-Feather sessions (BOFs) bring together devotees of many varied disciplines for discussions, announcements, mingling, and strategy sharing during evenings at the symposium. Schedule a BOF in advance or on-site by contacting the USENIX Conference Office, conference@usenix.org.

Work-progress-reports

These reports provide researchers with 10 minutes to speak on current work and receive valuable feedback. Present your interim results, novel approaches, or newly-completed work. Schedule your report in advance or on-site. The WIPS Coordinator is Greg A. Woods, woods@usenix.org.

For more information

Materials containing all details of the technical and tutorial programme, conference registration, hotel and airline discount and reservation information will be mailed in January of 1993. If you wish to receive the preregistration materials, please contact:

USENIX Conference Office
22672 Lambert St., Suite 613
El Toro, California
92630 U.S.A.
E-mail: conference@usenix.org
Phone: +1 714 588-8649
Fax: +1 714 588-9706

Important dates

Extended Abstracts Due:	December 4, 1992
Notifications to Authors:	December 16, 1992
Final Papers Due:	February 12, 1993

Call for Papers

ACM Multimedia '93, the First ACM International Conference on Multimedia, will be held in Anaheim, California August 1–6, 1993. The conference is co-located with SIGGRAPH '93, and is sponsored by the Association for Computing Machinery (SIGCHI, SIGCOMM, SIGGRAPH, SIGIR, SIGLINK, and SIGOIS) in cooperation with SIGAPP, SIGBIT, SIGBIO, SIGMOD, SIGOPS and the IEEE Communication and Computer Societies.

Topics

ACM Multimedia '93 will provide an international forum for papers, panels, courses, workshops, and exhibits focusing on the synergies between processing and communicating information represented in multiple media (multimedia). Research ideas, emerging technologies, engineering methodologies, prototype demonstrations, and experiences should be submitted for review. Technical areas for *Multimedia '93* include, but are not limited to:

- Applications and tools
- Collaboration environments
- Database and information systems
- Distributed systems
- Hardware and architectures
- Networking and communication
- Media integration and synchronization
- Image, video and audio compression techniques
- Operating system extensions
- Programming paradigms and environments
- Storage and I/O architectures
- User interfaces

Papers

High-quality technical papers on completed or in-progress research, innovative applications, and experience with multimedia systems are solicited. Where applicable, prototype demonstrations or videotape presentations are encouraged to supplement the talks.

Outstanding student and regular papers on different areas of multimedia will be given awards. Expanded versions of a small number of selected papers will be forwarded for possible publication in the *Communications of the ACM*, the *ACM Transactions on Information Systems*, and the new joint *IEEE/ACM Transactions on Networking*.

Panels

Panels are solicited that examine innovative, controversial, or otherwise provocative issues of interest. Proposals (not to exceed 3 pages, including biographical sketches of the panelists) will be evaluated based on technical quality, timeliness, appeal, and informativeness.

Tutorials

Proposals (at most 5 pages, including biographical sketches of instructors) for both 1/2- and 1-day tutorial courses are solicited. Evaluation of proposals will be based on expertise and experience of instructors, relevance of subject matter, and the use of multimedia technology in the presentation. Include needed audio-visual and computer equipment in your proposal.

Workshops

During the first two days of *Multimedia '93*, we would like to organize workshops on specific areas of multimedia research and technology. Evening sessions during the main conference (last three days) will be held to report the results of the workshops.

Please send proposals (at most 3 pages, including biographical sketches of organizers) for 1/2 day, 1 day or 1 1/2 day workshops, indicating if they have been held previously, limits on attendance, names of workshop chairs, keywords, title, length, a list and description of topics of interest, A-V needs, and (optional) sponsoring agencies.

Exhibition

Multimedia '93 offers a unique opportunity for vendors and researchers to exhibit and demonstrate multimedia products. There will be a section of the SIGGRAPH exhibition devoted solely to multimedia software and hardware. For exhibition information, contact Hall-Erickson, Inc., 150 Burlington Avenue, Clarendon Hills, IL 60514, +1 708-850-7779 • Exhibits.Multimedia93@siggraph.org.

Interactive media

At the conference, there will be a hands-on exhibit of interactive media projects. These projects will show state-of-the-art running systems in interactive media, and will include both desktop and networked systems. A two-page synopsis of the project along with an illustrative videotape (in VHS format) are solicited for review.

Electronic submission and publication

In addition to hard copy submission, papers may be submitted for review electronically via the Internet. Papers may use either plain ASCII text or *PostScript*. Multimedia papers should be submitted with a video walk-through, as well as a written two-page abstract. The videotape should be in VHS format, and six copies of the tape should be sent. The proceedings of the conference will be published electronically, as well as in hardcopy format. The format of the electronic version is currently being designed, and it is anticipated that this format will include support for rich media, and be accessible on both CD-ROM and electronic networks. Submitters must therefore be prepared to provide electronic versions of accepted material for publication.

Student participation

Papers with a student as the primary author will enter a student paper award competition; a maximum of two papers will be awarded a complimentary conference registration for the student author and a travel grant of up to \$500. A cover letter must identify the paper as a candidate for the student paper competition.

More information

An author's kit containing submission guidelines is available via FTP at siggraph.org. For conference information contact:

Dr. J. J. Garcia-Luna
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
Phone: +1 415-859-5647 Fax: +1 415 859-6028
E-mail: Chair.Multimedia93@siggraph.org

Submit printed papers/proposals (6 copies) to:

Prof. P. Venkat Rangan
Department of CSE,
9500 Gilman Drive, APM Building, Room 3016
University of California at San Diego
La Jolla, CA 92093
Phone: +1 619-534-5419 Fax: +1 619-534-7029
E-mail: Chair.Program.Multimedia93@siggraph.org

Send all electronic submissions to:

ElectronicSubmissions.Multimedia93@siggraph.org

Important dates

All submissions due: January 8, 1993
Notification of acceptance: March 1, 1993
Submissions in final form due: May 1, 1993

Call for Papers

Focus Prasun Dewan will be Guest Editor of a special issue of the journal *Computing Systems* to be published in 1993. The issue will be devoted to "Collaborative Computing Systems and Applications." Papers on all aspects of design, implementation, and experiences with these systems are solicited for the issue. The deadline for submissions is December 22, 1992; because of the holiday, papers submitted after this deadline will not be considered. Prospective authors should send five copies of their papers to:

Professor Prasun Dewan
 1398 Computer Science Building
 Department of Computer Sciences
 Purdue University
 West Lafayette, IN 47907-1398
 Phone: +1 317-494-6014
 E-mail: pd@cs.purdue.edu

Submissions Submissions should not have appeared in other archival publications prior to their submission. Papers developed from earlier conference, symposia and workshop presentations are welcome.

About the journal *Computing Systems* is a journal dedicated to the analysis and understanding of the theory, design, art, engineering and implementation of advanced computing systems, with an emphasis on systems inspired or influenced by the UNIX tradition. The journal's content includes coverage of topics in operating systems, architecture, networking, interfaces, programming languages, and sophisticated applications.

Computing Systems (ISSN 0895-6340) is a refereed, quarterly journal published by the University of California Press for the USENIX Association. USENIX is a professional and technical association of individuals and institutions concerned with breeding innovation in the UNIX tradition.

Now in its fifth year of publication, *Computing Systems* is regularly distributed to 4900 individual subscribers and over 600 institutional subscribers (libraries, research labs, etc.) around the world. Some special-topic issues are often distributed more widely.

The editor-in-chief of *Computing Systems* is Mike O'Dell of Bellcore. Gene Spafford of Purdue University is Associate Editor, and Peter Salus of the Sun User Group is the Managing Editor.

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now sixty-eight to choose from; ask for our free 1987-1992 index booklet). We want to hear from you. Contact us at:

ConneXions—The Interoperability Report
 480 San Antonio Road, Suite 100
 Mountain View, CA 94040-1219
 USA
 Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
 Fax: +1 415-949-1779
 E-mail: connexions@interop.com

Book Review

The latest entry into the Internet book realm is *Crossing the Internet Threshold: an instructional handbook* by Roy Tennant, John Ober, and Anne G. Lipow, with a foreword by Clifford Lynch. The publisher is Library Solutions Press and the price is \$40. The ISBN number is 1-882208-01-3. (I should note that I have a pre-release version of the handbook and not the final version.)

Two audiences

The book tries to serve two audiences at once. The first is the beginning Internet user and the other is Internet trainers as a training supplement. I fall into the second category. The handbook slightly leans towards librarians, but not enough that it hurts the readability for others. The handbook is a mixture of materials from lectures to overheads and one page summaries to exercises and checklists.

Concise and accurate

The handbook is one of the most non-threatening documents that I have seen on the Internet. The material is very concise without needless detail, but at the same time covers all the basics and is very accurate. One reason it is concise is that it describes the best documents, periodicals, and discussion groups for those wishing more than the basics. The material is so straight forward that I saw only one spot in the book where a user could get totally lost (the LISTSERV section) though on a second reading I realized that all the necessary steps were there and correct.

The one page summaries cover all kinds of related, but non-essential topics like BITNET, Gopher and Project Gutenberg. The exercises are well thought out. The overheads and such are of use to an Internet trainer who has not already designed similar materials.

On target

In summary, *Crossing the Internet Threshold* meets the needs of its target audiences. It is not a book for the experienced Internet user who does not train and it is not presented as such. My main complaint about the book is the price. The reason that has been explained to me is that this is the only title Library Solutions Press has and all of the overhead has to be covered by this book. I think it would make a useful addition to the book collection of both neophytes and Internet trainers alike.

—Billy Barron (billy@unt.edu)

A Letter to the Editor

Dear Ole,

The ongoing OSI discussion is entertaining to watch. ("OSI is (Still) a Good Idea," and follow-ups.) If OSI is a good idea, then stop discussions, join forces, and get it working! We are all waiting for it and its rich functionality. Complex software and complex protocols are no argument against OSI—just look what has happened in the PC industry during the last decade.

Regarding complexity: Take a look at Jon Postel's graph of the number of RFCs produced per month. What might be an explanation for the increase during the last few years, and what might be a consequence?

—Harald Hoffmann,
OSIconsult Kommunikationssysteme GmbH, Austria

CONNE~~X~~IONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNE~~X~~IONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Subscribe to CONNE~~X~~IONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNE~~X~~IONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNE~~X~~IONS

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

connexions@interop.com

Back issues available upon request \$15./each
Volume discounts available upon request

CONNE~~X~~IONS